TNO Defence, Security and Safety

Nederlandse Organisatie voor
toegepast-natuurwetenschappelijk
onderzoek / Netherlands Organisation
for Applied Scientific Research

ONGERUBRICEERD

TNO-report

TNO-DV1 2004 A253

A literature survey on Synthetic Aperture Sonar (SAS)

Oude Waalsdorperweg 63
PO Box 96864
2509 JG The Hague
The Netherlands

www.tno.nl

Phone +31 070 374 00 00
Fax +31 070 328 09 61
Info-DenV@tno.nl

Date	April 2005
Author(s)	J.C. Sabel, S.P. Beerens and J.Groen
Classification	Ongerubriceerd
Affiliation	DMKM/WCS/OSWS
Project officer	LTZE 2OC J.P. Gillesse
Classification date	21-01-2005
Title	Ongerubriceerd
Managementuittreksel	Ongerubriceerd
Abstract	Ongerubriceerd
Report text	Ongerubriceerd
Appendices	-
Contract no	A99KM757
Sponsor	HWO-KM
Affiliation	KM
Project name	OnSonar
Project number	015.33053
Copy no	5
No of copies	17
No of pages	37 (incl. Appendices, excl. RDP & distribution list)
No of appendices	-

Netherlands Organisation for Applied Scientific Research (TNO)

A literature survey on Synthetic Aperture Sonar (SAS)

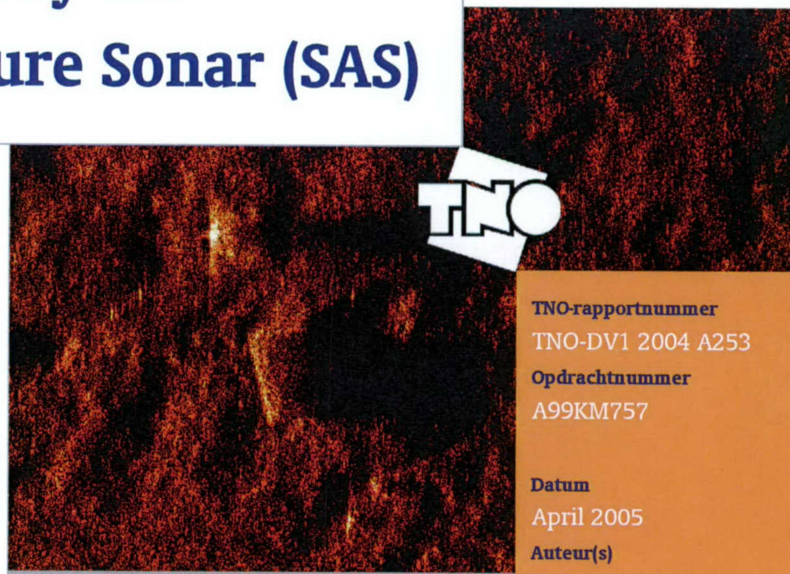

TNO-rapportnummer
TNO-DV1 2004 A253

Opdrachtnummer
A99KM757

Datum
April 2005

Auteur(s)
Dr.ir. J.C. Sabel,
Dr. S.P. Beerens
Ir. J.Groen

Rubricering rapport
Ongerubriceerd

Probleemstelling
In het kader van het project SAS (Synthetische Apertuur Sonar) is een literatuurstudie uitgevoerd naar bestaande methodes voor SAS processing en bewegingscompensatie. Het doel ervan is om bekend te raken met het onderwerp SAS. Dit rapport dient als samenvatting van de studie en als naslagwerk voor de rest van het project.

Beschrijving van de werkzaamheden
Er is gezocht naar wetenschappelijke publicaties en boeken over SAS. Ook het onderwerp Synthetische Apertuur Radar (SAR) is bekeken, vanwege de overeenkomsten tussen SAS en SAR processing. Over SAR bestaat bij de Radargroep van TNO-FEL veel kennis, waarvan expliciet gebruik wordt gemaakt. De gevonden publicaties zijn gelezen en gerangschikt naar onderwerp en relevantie.

Resultaten en conclusies
We hebben (nog) geen standaardwerk kunnen vinden op het gebied van SAS. De gevonden artikelen behandelen meestal een deelaspect. Voor een basisoverzicht moet teruggegrepen worden op SAR literatuur. Voor het project zijn diverse gevonden artikelen over SAS goed te gebruiken, juist omdat de specifieke problemen daarin aan de orde komen. Veel literatuur is recent en afkomstig uit proceedings van conferenties.

ONGERUBRICEERD

TNO Defensie en Veiligheid

Oude Waalsdorperweg 63
2597 AK 's-Gravenhage
Postbus 96864
2509 JG 's-Gravenhage

www.tno.nl
Info-DenV@tno.nl

T 070 374 00 00
F 070 328 09 61

PROGRAMMA	**PROJECT**
Programmabegeleider -	Projectbegeleider LTZE 2OC J.P. Gillesse, DMKM/WCS/OSWS
Programmaleider -	Projectleider Dr.ir. J.C. Sabel, TNO Defensie en Veiligheid
Programmatitel Sonar	Projecttitel SAS
Programmanummer V 036	Projectnummer 015.33053
Programmaplanning Start 01-01-2001 Gereed 31-12-2004	Projectplanning Start 01-03-1999 Gereed 31-12-2002
Frequentie van overleg -	Projectteam Dr.ir. J.C. Sabel, Dr. S.P. Beerens Ir. J.Groen M.E.G.D. Colin

Contents

Abbreviations

AFNCP/SCM/WG3	Anglo-French-Netherlands Collaboration Programme/ Subcommittee M (Mine warfare)/ Working Group 3
ASW	Anti-submarine warfare
CPU	Central processing unit
CW	Continuous wave (signal)
DERA	Defence Evaluation and Research Agency
DPC	Displaced phase centres (autofocus algorithm)
FFT	Fast Fourier transform
FM	Frequency modulated (signal)/modulation
GESMA	Groupe d'Études Sous-Marines de l'Atlantique
MCM	Mine countermeasures
NAT	New array technology
PGA	Phase gradient autofocus (autofocus algorithm)
PRF	Pulse repetition frequency
PTR	Point target response
SAR	Synthetic aperture radar
SAS	Synthetic aperture sonar
SNR	Signal to noise ratio
TNO	Netherlands Organisation for Applied Scientific Research
UUV	Unmanned underwater vehicle

1. Introduction

Modern mine hunting requires high resolution sonar systems for countering low-target-strength (or *stealth*) mines. Wide-band techniques enable very high range resolution, in the order of centimeters. However, conventional sonars are limited in *azimuth resolution* by their real acoustic aperture (physical array length); increasing this resolution would require a prohibitively large receiver array. Furthermore, this resolution is range-dependent. In other words, a real aperture has a fixed angular beam width. Synthetic aperture sonar (SAS) is generally recognized as the technique that can be used to solve this limitation.

The basic approach of SAS is the formation of a sonar image by coherently integrating many successive pings from a moving platform. This yields a synthetic array that is much longer than the physical array. The synthetic array length increases with distance as a consequence of the transmitter and receiver beam patterns, thereby providing an absolute resolution that is (at least theoretically) range independent. In practice, the resolution at long ranges will be limited by several factors, such as the detection range of the sonar and decorrelation of the signal due to propagation effects.

SAS is an attractive way of achieving high-resolution sonar imaging with unmanned underwater vehicles (UUVs) due to the limited array length requirements. These platforms are being increasingly applied in mine hunting, to enable low risk stand-off surveillance.

Many aspects need to be considered for obtaining good images with SAS, such as phase errors due to platform movement and propagation effects. So-called autofocus techniques are required to correct for these effects. Another important aspect in SAS processing is the reduction of computational load.

SAS has not yet fully matured, unlike its radar equivalent, synthetic aperture radar (SAR). SAS is still in the research and development phase, whereas SAR has been operational for many years. As a consequence, books on SAR are available, e.g. Curlander [1] and Carrara and colleagues [2] or for an overview, Otten [3]. However, there seem to be none on SAS and most SAS publications are conference or journal papers, dealing with details rather than basics.

History
The first good review paper in which the important parameters of SAS are identified and quantified was published in 1975 by Cutrona [4]. We know that SAS is much more complicated than SAR [5, [6], mainly due to the limited speed of sound compared to the speed of light. This means that the pulse-repetition frequency (PRF) will generally be very low in SAS, which hinders straightforward synthetic aperture formation. Long integration times are required in order to collect

a sufficient number of pings. Much may change over this extended period, such as; target bearing, target range, the medium, the platform speed/orientation, etc. This makes SAS processing more complicated and this is the main reason why SAS has become popular so many years after SAR. However, in the 1990s many studies were published [7-[9]. Research programmes on SAS are currently running in several countries, such as the USA, New Zealand, Sweden, France, Great Britain, Australia and the Netherlands, and in institutes like Saclantcen.

Problems

Many problems need to be solved compared to SAR before SAS is a feasible alternative. It is also possible that completely new concepts not linked to SAR have to be developed in order to solve these problems. Poor sampling due to low pulse-repetition frequencies [10] is a problem with SAS that does not occur in SAR. This is due to the relatively low speed of sound in water compared to electromagnetic waves, in combination with the requirement of sampling with a spacing of half an array length. The same effect poses a restriction on the area coverage rate –the product of swath width and platform speed– that can be achieved.

Another fundamental problem in synthetic aperture formation is the uncertainty in platform position [12, [11]. This uncertainty leads to phase errors on the elements in the array, which can have a dramatic effect on the received beam pattern and the resulting image. We know from beam-forming theory that the position of the elements should have an accuracy of least ¼ of an acoustic wavelength for meaningful processing. Even an accuracy of $^{1}/_{10}$ of an acoustic wavelength is required for it to perform as well as a real aperture array [36]. Extremely accurate knowledge of the sonar position is required since mine-hunting frequencies are high, and thus wavelengths are short (centimetres)..This is unrealistic with the present accuracy of positioning sensors. However, *autofocussing* may help here and this is a current 'buzzword' in SAS [15-[23]. Autofocussing uses the acoustic signals themselves to estimate and compensate for the array position and orientation.

There are additional computational difficulties for application of SAS in modern mine-hunting sonars. These are mainly due to the large bandwidth-time products that are currently available. The advantages of large bandwidth pulses are increased range resolution and better performance against reverberation, which leads to enhanced classification and detection performance. However, the processing of such large bandwidth-time products is computationally expensive and therefore many efforts are being made to reduce the computational load [24-[29], especially by those who exploit real-time demonstrator types of SAS systems. Meaningful processing of the ordinary aperture is already difficult without special measures [9], let alone the processing of a synthetic aperture. The computational problem of beam-forming with wideband pulses is also an important topic in modern anti-submarine warfare (ASW) sonars. Some expertise can be borrowed from this field. Finally, SAS beam-forming is near-field processing by definition, which means

that the change of target range must be taken into account. Targets at different ranges have a different wave-front curvature, which must be taken into account in the processing stages.

Related problems
ASW is another application area of SAS, although outside the scope of the present research. Modern active ASW arrays are relatively short and, in order to be able to use them for passive detection at low frequencies, synthetic aperture formation is a serious option. These low frequencies become of increasing interest, as submarines with air-independent propulsion are becoming more widespread. Although 'passive SAS' is little less complicated than 'active SAS' that we want to use in mine countermeasures (MCM), it still meets a lot of similar problems. Several papers have been published in this field that are of interest for this study, mainly by Stergiopoulos [37-[41] and Williams [42,[43].

Outline of this report
The remainder of this report is organised as follows: the general principles of SAR are explained in Chapter 2. The principles of SAS are basically the same; however, there are some additional problems and these are treated in Chapters 3 and 4.
A discussion on possibilities, problems, challenges and the future for SAS are discussed in Chapter 5.

2. Principles of SAR

We first discuss SAR processing, which could be considered as the cradle of SAS, since SAR and SAS processing are very similar. We quote Otten [3] and the references therein in this chapter. The chapter is mostly intended as an introduction to SAR processing for those who are acquainted with radar or sonar, but not so much with SAR or SAS.

2.1 Basic principles of SAR and SAR processing

Since the beginning of SAR, many different ways of processing the raw radar data into SAR images have been devised. In this chapter, the fundamentals of SAR processing are explained, and an overview of all different algorithms is given, outlining the reasons for their existence.

SAR employs a side-looking radar carried by an aircraft or satellite that images the surface below. The synthetic aperture is formed by collecting multiple pulses along the flight path of the platform. Synthetic aperture processing combines the principles of *pulse compression*. Pulse compression (in sonar terminology this is generally called matched filtering) is used in many radar and sonar applications to create the effect of a short pulse, for high range resolution, by using a long modulated pulse that carries more energy. Effectively a short pulse is obtained by matched-filtering the received pulse, resulting in a high signal-to-noise ratio (SNR).

The azimuth (cross-range) resolution of a side-looking radar is determined by the antenna beam width, and therefore by the real size of the antenna aperture. In conventional (real aperture) radar this aperture should be as large as possible. In SAR, however, the real aperture is synthetically extended, by *coherent* integration of the received signal over many pulses along the direction of motion of the radar. In this way, a narrow synthetic beam is created, which provides the high azimuth resolution that is typical of SAR.

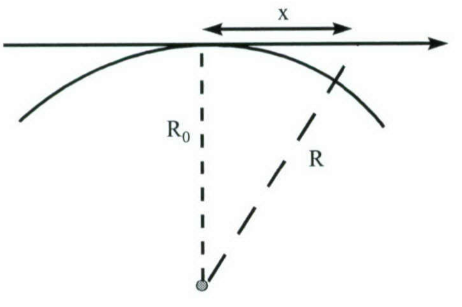

$$R = \sqrt{R_0^2 + x^2} \approx R_0 + x^2 / 2R$$

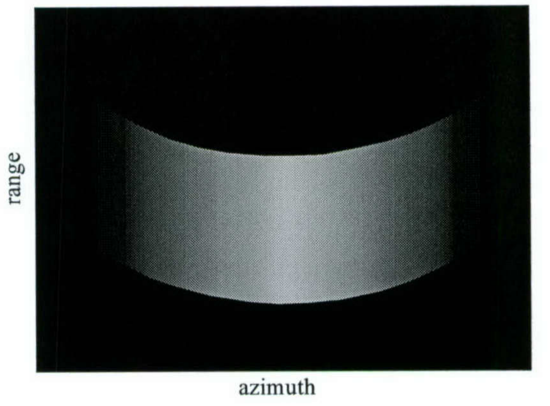

Figure 2.1: Range variation to a point

Figure 2.2: Typical shape of the point target response (PTR) before range and azimuth compression

Mathematically, synthetic aperture formation is very similar to pulse compression, and is generally referred to as *azimuth compression*. In fact, just like the linear FM pulse that is often used for pulse compression, the azimuth response to a point target is also well approximated by an FM signal. This response is Doppler shifted[1] by the relative motion between the radar and the target, and varies from a positive Doppler, when the radar is still approaching the target, to a negative Doppler, when the radar has passed the target. Around the point of closest approach, the Doppler shift varies almost linearly, which explains the linear frequency modulation in azimuth. This can also be derived easily with the help of Figure 2.1, by considering the range (R) from the radar to a point when the radar is passing by, as a function of position (x):

$$R = \sqrt{R_0^2 + x^2} \approx R_0 + \frac{x^2}{2R_0} \tag{2-1}$$

The quadratic term in the range leads to a quadratic phase term in the signal, since the phase φ is proportional to the travel time of the wave. A quadratic phase term corresponds to a linear frequency term.

$$\varphi = \frac{4\pi}{\lambda} R \approx \varphi_0 + \frac{2\pi x^2}{\lambda R_0} \tag{2-2}$$

and, with $x = Vt$, where V is the platform speed:

$$f = \frac{1}{2\pi} \cdot \frac{d\varphi}{dt} = \frac{2Vx}{\lambda R_0} = \frac{2V^2 t}{\lambda R_0} \tag{2-3}$$

---

[1] Note that in radar the term "Doppler" is used for range migration between different pings, whereas in sonar it is used for range migration (=speed) within a ping.

Since the range and azimuth dimensions of the signal are in fact only dimensions of time, they are also referred to as fast time (nano- to microseconds) and slow time (milliseconds to seconds). The factor t in (2-3) refers to slow time.

Thus, SAR processing is basically a matched filtering of received data with a 2-D frequency modulated signal. Although simple in principle, SAR processing in practice is a complicated task as explained below.

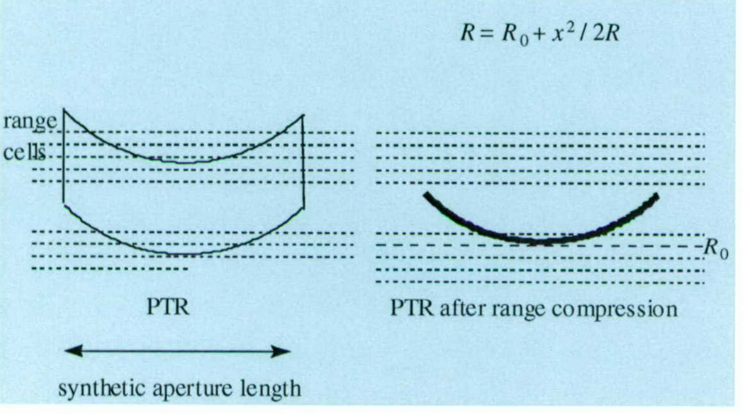

Figure 2.3: Schematic representation of the PTR and the range migration effect

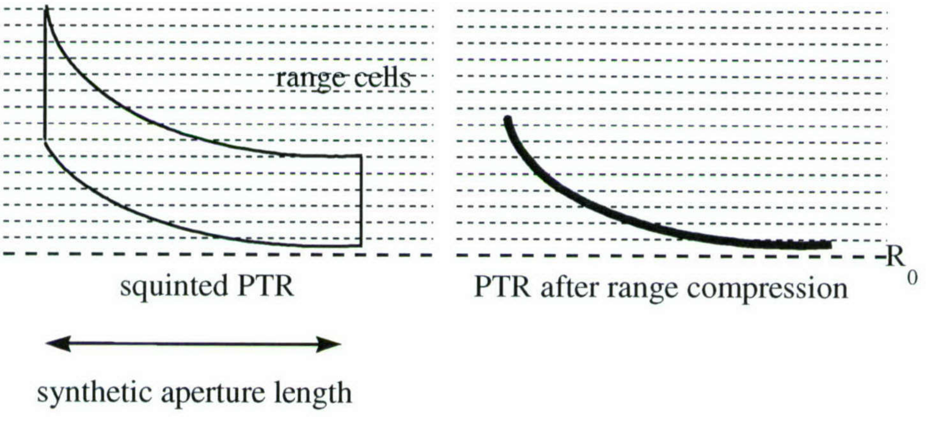

Figure 2.4: Point target response (PTR) for a squinted SAR (looking slightly forward)

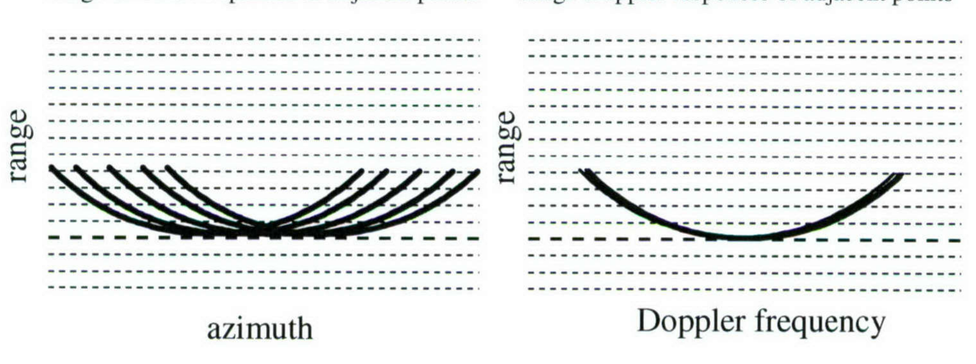

Figure 2.5: Range-azimuth vs. range-Doppler migration

2.2 SAR processing algorithms

In the previous section we showed that SAR processing is a coherent summation of the target response in two dimensions. Mathematically this is a 2-D matched filtering, matched to the response of a point target. Range modulation is imposed by the system, and is often a linear FM sweep. Azimuth modulation is also an FM, which is a consequence of the data collection geometry. The amplitude is determined by the uncompressed pulse shape in range, and by the antenna beam pattern in azimuth. The typical point target response shape is shown in Figure 2.2. We assume here that the pulse amplitude is not modulated, and thus has a rectangular pulse shape. Obviously, the range extent has clear boundaries, while the azimuth extent is fuzzier due to the transmitter beam pattern. The main lobe of the azimuth response is shown here, not the side lobes. The main lobe is the beam that is used for standard SAR processing (in spotlight SAR, the target is followed and other beams are also processed).

In the simplest case this operation, which is in fact a 2-D convolution, can be separated into two 1-D convolutions. These can be carried out quite efficiently using fast convolution, i.e., convolution via frequency-domain multiplication. A special case is when the synthetic aperture is so short that even the azimuth Doppler modulation can be neglected, so that mere integration is sufficient. This is called *unfocused* SAR and it applies only to low-resolution SAR systems.

More often, separation into two 1-D convolutions is not possible. Although the convolution can be separated into range and azimuth compression, the latter is then no longer 1-D, due to the fact that the range-compressed point-target response migrates over several range cells during the synthetic aperture formation time. Therefore the azimuth compression becomes 2-D, and the different approaches to cope with this lead to different types of algorithm. Range migration is illustrated in Figure 2.1, showing how the range-compressed response migrates over several range cells. The effect is worse when the SAR is squinted, i.e., when the antenna is not pointing in a direction perpendicular to the direction of motion of the radar. In that case, the target looks more like that in Figure 2.2. Note that the amount of range migration increases, and also the azimuth length within one range cell becomes very short: the azimuth signal within one cell then begins to lose its linear FM character.

Another complication in SAR processing is the fact that the required 2-D convolution kernel is not fixed; it may vary over both range and azimuth. The range (R_0) dependence is systematic and is apparent from eqn. (2-3). This means that in practice, the convolution must be separated into convolutions over smaller amounts of data, which can be processed with fixed parameters. This separation into smaller parts can be a significant performance bottleneck, depending on the type of algorithm.

One solution to 2-D azimuth convolution is not to split range and azimuth compression in the first place, but to perform 2-D convolution with the complete range/azimuth response (as shown in the left half of Figure 2.3). This type of processing can be done using the full, exact point target response (i.e., transfer function), and is also called exact transfer function (ETF) processing. This type of algorithm can be quite efficient if there is only a small spatial variance and the 2-D templates are not too large; however, memory requirements are high.

Another very straightforward approach is to perform the azimuth integration along the range migration curve; this requires many interpolations, and a different curve must be followed for every image pixel This approach is extremely computationally intensive and is no longer used in practice.

One of the first solutions to this problem that has become very popular is the so-called Range-Doppler algorithm. It makes use of the fact that the time and frequency of the azimuth signal from a point target are uniquely coupled, so that range migration correction can also be performed in the transformed range-Doppler domain instead of the range-azimuth domain. The advantage of this is that the migration curve in the range-Doppler domain is the same for many points on the ground, so that one migration correction can be performed for a whole block of image pixels in this domain (Figure 2.3). This algorithm became known in the early 1980s, and has become widespread due to its efficiency and simplicity. There is one major drawback: since the relation between slow time and Doppler frequency is not perfect, the image will exhibit degradation when the amount of range migration becomes too large. This effect can be reduced by so-called secondary range compression. However, this is still insufficient for very highly squinted data.

There are also other ways of performing 2-D azimuth convolution: often the azimuth extent of the reference will be long, while the range extent is short. The azimuth dimension can be convoluted via the frequency domain, and the range dimension in the time domain. This is called the hybrid correlation approach and the advantage of this approach is that no approximations need to be made, so that the algorithm is very accurate; the main disadvantage is that the efficiency drops when the range migration increases.

Two more algorithm types need to be mentioned here: the so-called 'ω-k', or wavenumber domain algorithms, and the fairly recent 'chirp scaling' algorithms. The former was originally an adaptation of seismic processing algorithms to SAR. It requires a transformation, called Stolt interpolation, in the 2-D frequency (or wavenumber) domain, which is also its most critical stage. One of its attractive properties is that the range-dependence of the SAR transfer function is inherently included. There are several approximations to this algorithm, which vary in complexity and accuracy.

The latter algorithm, chirp scaling, makes explicit use of the linear FM (or 'chirp') property of the transmitted pulse, and is therefore not generally applicable. It applies a chirp rate alteration by phase multiplications to range-Doppler data leading to shifting of the range correlation maxima, which can then be made to compensate for the range migration. This algorithm has also become popular, because it compensates migration without the need for interpolations, and it has also been further improved recently.

Another factor leading to changes in algorithms over the last few years is the fact that SAR interferometry is becoming a more established application of SAR. Although formerly the *intensity* was the most or only interesting property of an image pixel, the *phase* of image pixels has become essential in interferometric applications. Therefore, SAR processors now need to be 'phase–preserving' and most modern algorithms satisfy this condition.

2.3 Motion Compensation

Unfortunately, using a smart compression algorithm is not the end of the story. Many additional measures must be taken especially in airborne SAR to compensate for undesired and non-stationary motion of the aircraft. First of all, the relevant motions must be known. This is achieved either by direct measurement from motion sensors or by parameter estimation using the radar data. Motion is measured mainly by acceleration (inertial) sensors on the aircraft, nowadays often combined with GPS position data to improve the long-term accuracy. Estimating motion from the radar data is the domain of autofocus algorithms. In fact, these algorithms actually estimate the required focusing parameters, which are the result of the imaging geometry and the unknown motions. Without going into detail, two main algorithms are mentioned here: map drift autofocus, which relies on measuring the shift between multiple looks, and phase gradient autofocus, which estimates the azimuth phase history derivative from averaging over range of scatterers which are first aligned in azimuth.

Motion compensation, i.e., applying motion measurements or estimates to processing, may take on many different forms and shapes, depending also on the type of SAR system. It is obvious that, for instance, an airborne SAR using an inertially stabilised antenna platform does not require the same compensation as a system in which the antenna is rigidly fixed to the airframe. It is not possible to treat every existing compensation scheme, but Otten [3] works out an example of the Dutch PHARUS system, which is a phased-array airborne SAR.

2.4 Advanced techniques

New modes of SAR operation have come into existence in addition to SAR strip-mapping modes. Interferometry has already been mentioned. Some other modes are 'spotlight' and 'ScanSAR' modes. Spotlight is actually not new, but has been mainly applied in military systems: in spotlight mode, the antenna beam is trained on one spot on the ground, to increase the integration time, and hence, the resolution. So-called 'polar formatting' is the type of algorithm that is often applied in this case. A relatively new mode is ScanSAR, in which the antenna beam is quickly scanned in the elevation direction, to obtain larger range coverage than that obtained in normal strip-mapping mode. This type of mode can only be carried out with fast electronic beam steering, i.e., with phased-array antennas. ScanSAR is used operationally, as for instance in the Canadian RADARSAT satellite. One of the particular difficulties in ScanSAR processing is the fact that the subswaths making up the whole image have all been illuminated with different beam positions. A high degree of processing accuracy is required to obtain radiometric consistency over the whole SAR image.

2.5 Discussion

SAR processing can be described as performing 2-D convolution with a complex function. The main complications in this process are:

- size of the function, and amount of data;
- spatial variation of the function;
- range migration;
- motion compensation.

Different SAR algorithms and processors have different ways of coping with these complications. Well-known algorithms are the widespread range-Doppler algorithm, later improved with secondary range compression, wavenumber domain algorithms, and more recently chirp-scaling algorithms, which exploit the linear FM property of the transmitted pulse. Advanced SAR applications, like interferometry, have added the condition of phase preservation to the SAR processor requirements. There is no ideal algorithm; the choice depends on the SAR characteristics, and the requirements for speed and accuracy.

Motion compensation in the processing tends to be very SAR-system dependent, since different airborne systems have different ways of dealing with motion. Some use mechanical antenna stabilisation, while others compensate in the processing stages. Moreover, motion compensation requirements can be quite different for short-range or long-range systems, and for small or large airborne platforms. Therefore, a good understanding of the SAR system and the applications is required.

Currently, much research effort is being devoted to advanced processing techniques related to SAR, such as SAR/ISAR imaging. New techniques for moving target indication, and moving target imaging are also receiving attention [3].

2.6 Relevance to SAS

SAR and SAS processing are similar in many respects and much of the above-mentioned methods and problems also apply to SAS. However, there are also differences.

The main overlap occurs in SAR processing. Many of the SAR algorithms and processors can be applied in SAS and, as there is no ideal algorithm, many algorithms must be tested. Most SAR algorithms have been designed for computational speed, rather than for accuracy. We aim primarily at accuracy for SAS; computational speed has a lower priority.

An important difference with SAR is that in sonar we usually work with element-level signals, allowing multi-beam beam-forming in post-processing, whereas in radar there is usually only one beam per pulse. This opens up a wider range of

processing options. For instance, it allows spotlight processing for the entire observed area in SAS where that would only have been possible for a relatively small spot in spotlight SAR.

Another difference occurs in motion compensation. Sonar platforms move differently from radar platforms and this will influence the correction for these movements by motion sensors and/or by autofocussing. Autofocussing in SAS will generally be more difficult because the variations are larger and the synthetic aperture covers a longer time span, and because the overlap of consecutive array positions will generally be less due to the limited PRF.

Finally, environmental influences are more important in sonar than in radar. Propagation effects tend to reduce signal coherence and a lack of features on the average sea bottom may cause problems in autofocussing methods that rely on contrast or correlation in the image.

Summarising, some techniques are very similar in SAR and SAS and can be exchanged, but unfortunately many aspects are quite different. There is still work to be done here. The following chapters mention the main points that deserve attention.

3. Real aperture processing of wideband MCM arrays

SAS processing differs from SAR processing mainly in two aspects: first the processing of the real aperture in SAS is more laborious, and second the actual formation of the synthetic aperture is more complicated than SAR. These two processes are rather different and therefore they are treated separately in the two following chapters starting with real aperture array processing.

3.1 Introduction to real aperture beam-forming of SAS arrays

Before we can start with the SAS process we must first carry out beam-forming on the real aperture, which is generally done in hardware in SAR. Although sonar array processing seems standard and is quite well understood [30] this topic never seems to vanish from sonar research. The reason is the continuous rapid advance of modern sonar systems and their associated computational power. Sonar arrays become increasingly longer, relative bandwidth increases rapidly and the increased computational power makes it tempting to apply new algorithms, such as *array calibration* and/or *adaptive beam-formers* [31]. Adaptive beam-forming is especially computationally expensive, but has proved successful lately in many sonar applications, especially in environments with strong directional noise. For example, in submarine flank-arrays, these beam-formers are very promising.

3.2 Array calibration

A hot topic in modern beam-forming is *array calibration*. We know that defective hydrophones can reduce the quality of the beam pattern dramatically. Only one bad element in a 100-element array can give an increase in the side-lobe levels of the beam pattern of several dB. If we want to apply array shading (Section 3.3) calibration is essential. A defective element results in a zero shading coefficient, which may ruin the effect of the overall shading, especially if this element is close to the centre of the array. Even if only a few elements are bad, the effect on the side lobes may be adverse. Calibration can restore bad elements and reinforce the shading procedure. If we do not calibrate, it is better to skip the shading for defective elements.

Calibration is a special issue for adaptive beam-forming (Section 3.4), since adaptive beam-formers are very sensitive to phase errors. Phase noise due to bad electronics may destroy all favourable properties of an adaptive beam-former [9].

It should be noted that it is much better to avoid all of the above problems. We need to make sure that the array is of very high quality and well tested and calibrated before going to sea. The application of calibration methods is merely a

tool to repair sonar performance to some extent, but the best results are always obtained with a good array rather than with a bad array and fancy processing. It may be strange that such a remark should appear in a scientific report, but all too often sonar processing engineers are confronted with bad data and then nothing can be done, not even with the most cleverly designed algorithms.

The most simple and most often used solution is to remove all bad channels. For the beam pattern, a hole in the array is not as bad as a bad element. The beam-forming operation has an interpolating effect on the hole. However, two main problems occur in this treatment: how to detect and classify a bad element? Most of the time elements are only temporarily bad. How bad should they be before they are removed?

Instead of removing them, repair seems much more attractive. Amplitude calibration is not difficult, and for each element a gain factor can be applied. Note that for wideband these are frequency dependent. Phase calibration however, is much more complicated. At DERA an algorithm is applied (ROSCO [32]) that retrieves the phase from interpolation and extrapolation with non-defective neighbours. This algorithm is robust in the sense that it always improves sonar performance, but interpolation still does not retrieve the lost phase information very accurately. *Auto calibration* methods could retrieve this phase information. Liu suggests in his thesis [9] the use of the acoustic channel response function. The echo returns of a target of opportunity (bottom) can be used for this purpose. In the NAT II research programme on low frequency active sonar it is suggested to transmit low-level calibration pulses in between the real sonar detection pulses. These auto-calibration methods seem successful [33]. Which of the calibration methods is the most suitable for application in SAS depends on the quality of the receive array and on the choice of the beam-former. This will be the subject of further study.

3.3 Array shading

In a sonar beam pattern, side lobes should have a low level to prevent reverberation and directional noise from different bearings leaking into the steering direction. In order to reduce side-lobe levels we recommend applying *array shading*. Shading controls the side lobes at the expense of a wider main lobe. Since in SAS processing the azimuth resolution is determined by SAS beam-forming rather than by the real aperture beam-former, it is less important to have a narrow main lobe than to have low side lobes. Therefore, shading (e.g., Dolph-Chebyshev or standard Hamming [30]) is even more strongly recommended in SAS than in other sonar processing methods.

3.4 Wideband beam-forming: time or frequency domain

For mine-hunting applications, typical array lengths are 100 elements in order to obtain sufficiently narrow beams and array gain. This number is quite typical in many sonar and radar applications and will not scare computational engineers. The wide bandwidth however, (which is extremely useful; see section 3.2.2), in combination with long pulses, does provide a large computational burden. Large BT (bandwidth-time) products require long FFTs (fast Fourier transforms), which are relatively slow and consume lots of memory.

Relative bandwidths are small in SAR and are generally expressed using the quality factor $Q = f / B$, which is small if the relative bandwidth is large. This means that engineers can get away with a so-called 'monochromatic' approximation in SAR with high Q, i.e., the FM signal is treated as if it was a CW signal. This gives a large computational advantage over SAS. We can have Q factors as low as 2-3 in SAS, which means that the upper and lower frequencies in the band differ substantially. In this case 'monochromatic' (or 'monotonal') approximations fail, since many relevant physical processes are frequency dependent, such as:

- phase differences between elements;
- element sensitivity;
- transmitter directivity index;
- receiver directivity index;
- absorption loss;
- the ambient noise and reverberation background;
- Doppler shifts.

These frequency dependencies should be accounted for in beam-forming. This provides a large computational burden and much effort is paid to reduce this burden. Quasi-narrowband methods have been suggested [24 - [26]. These are all time-domain beam-formers.

It is true that time-domain beam-formers are generally faster for wideband sonar [28], but they cannot deal with all of the above frequency-dependent processes. As long as real-time aspects have no priority in the research and development phase, it is not wise to make concessions to accuracy in order to reduce computational load. Thus, our first aim was to develop an accurate beam-former, preferably in the frequency domain, such that all above frequency-dependent effects can be accounted for. Another advantage of working in the frequency domain is that matched filtering in this domain is more efficient. It is faster in the frequency domain and since pre-whitening can be applied more easily, it is also more effective. If real-time aspects are considered in the future, it is hoped that hardware (CPU and memory) will have improved sufficiently to accommodate this processing.

3.4.1 Adaptive beam-forming

Adaptive beam-forming is becoming very popular nowadays instead of standard delay and sum beam-forming. In a standard beam-former all signals are added coherently to maximise the signal gain. It is expected that noise adds up incoherently, such that noise gain is minimal. In practice however, this is rarely the case. Noise is always partially coherent, especially when directional noise (shipping or own platform) is involved. The approach of an adaptive beam-former is to maximise the SNR rather than the signal, which is optimised in a standard beam-former. It is usually much more efficient to minimise noise, than to maximise the signal.

The optimum beam-former in this case is a complicated weighted sum of all element signals. The non-linear weighting coefficients depend on bearing, frequency and time. The latter complicates the matter as we need to adjust the beam-forming coefficients all the time to adapt to the prevailing signal and noise conditions. This makes the method extremely computationally intensive. Another disadvantage is that since the method is non-linear, it is very sensitive to small phase errors. The method tends to fail if the phases are erratic due to element positioning or bad electronics,. The reason for this is that cancelling of directional noise requires proper zero-steering and low side lobes and thus accurate phase information. However, despite these problems, adaptive beam-forming can yield substantial benefits over standard beam-forming in many environments.

In Liu's thesis [9] adaptive beam-forming was tested for an MCM array. Under the motto "*try simple things first*" he applied the 'MUSIC' algorithm, which is a standard in adaptive beam-forming. The algorithm was developed for 'narrowband passive' and was adapted for 'wideband active'. In wideband active sonar we encounter the problem that the time (=range) dependence of the beam-forming coefficients becomes very high. They need to be adapted very rapidly and in such a short time it is not possible to cope with frequency dependency (for a proper short time FFT you need more than just a few samples). It is better to carry out beam-forming in the time domain in that case [34].

The results of Liu are remarkable. After array calibration, he found that adaptive beam-forming yields a higher signal-to–background ratio (noise, reverberation and side-lobes) than standard beam-forming by about 20 dB (without calibration, results are worse than the standard by about the same amount). Clearly, adaptive beam-forming combined with calibration helps in a hostile MCM environment with much directional noise and reverberation. However, a big disappointment followed.

In the next processing step, SAS beam-forming, it was found that the results of adaptive beam-forming are so fragile that coherent SAS beam-forming was no longer possible. This is obviously related to the earlier-mentioned sensitivity to array positioning. Non-coherent SAS (a multi-ping method which adds amplitudes incoherently rather than the complex numbers including corrected phases

coherently) is still a possibility with the adaptive beam-forming results. Non-coherent SAS helps classification, but does not improve detection as the SNR does not improve.

The overall conclusion was that standard beam-forming with coherent SAS outperforms adaptive beam-forming with non-coherent SAS and this again demonstrates the power of coherent SAS.

It is possible that Liu's results can be improved. Liu ignored the important point that adaptive beam-forming is <u>not</u> necessarily a fully coherent summation. Hydrophone signals are added in such a way that the SNR is maximised, whereas in ordinary beam-forming hydrophone signals are added coherently (in phase) to maximise signal level. Therefore, the phase of a signal after adaptive beam-forming is 'unpredictable', and since adaptive beam-forming coefficients vary from ping to ping, phase deviations vary over time. It is not strange that signals with unpredictable time-varying phases cannot be added coherently by ordinary SAS algorithms. Something very clever must be discovered before adaptive beam-forming in combination with coherent SAS becomes feasible. This could however be a significant step forward, and needs proper consideration.

4. Synthetic aperture sonar processing

Although the main line of processing as outlined in Chapter 2 is followed, SAS processing is more complicated than SAR processing. The main cause of this is the limited speed of sound compared to the speed of light. This means that the pulse-repetition frequency will generally be low in SAS (typically around 1-5 Hz), which hinders straightforward synthetic aperture formation. Long integration times are required in order to collect sufficient pulses. Many things may change in this period of time, such as target range and aspect, medium/propagation conditions, platform speed/orientation, etc.

Other complications are caused by the more complicated medium. The underwater environment is hostile, with high propagation loss and all kinds of unwanted scattering lead to high reverberation levels. Several tasks have been defined for theoretical and environmental studies (Task 2.1 – 2.9) in order to deal with these problems within the SAS project. This chapter treats the relevant literature on these tasks in nine separate sections. Clearly not all problems are equally substantial and consequently not all problems get equal attention.

These tasks deal with environmental effects that influence SAS performance and with special SAS modes, which may be more sensitive to environmental degradation than standard SAS methods.

1. Sea floor reverberation
2. Water column influences
3. Surface reverberation
4. Multi-path in shallow water
5. Array motion compensation by autofocussing
6. Interferometry
7. Increased coverage rate
8. Buried mine detection
9. Shadow and echo classification.

4.1 Sea floor reverberation

Whereas sea floor reverberation is usually a disturbing factor in other sonar application areas, it is simply the background against which to detect targets in high-resolution sonar. The sonar image shows the sea floor in reasonable detail and in that sense corresponds to what we actually want to visualise. Without sea floor reverberation there would be no target shadow, which is a helpful feature in mine detection and classification tasks.

Reverberation is exploited in yet another respect. Displaced phase centre autofocussing (DPCA), for instance, relies on correlation of received signals from consecutive pings. If there is no reverberation, the correlation cannot be calculated properly. Thus, the reverberation will influence the autofocussing process and detailed reverberation properties may influence the robustness of autofocus techniques. This makes the study of sea floor reverberation, and especially its coherence properties, relevant for SAS. No specific publications have been found on this topic.

4.2 Water column influences

Non-homogeneities in the water column (e.g., hot or fresh water on top of other water masses) make the sound paths bend. Moreover, if these non-homogeneities are time dependent, as is the case when turbulent eddies occur, sound can have bent paths leading to phase distortions and loss of coherence in the signals. This can have serious consequences: beam-former and matched filter gain will decrease and, more importantly, autofocussing and related SAS processing may suffer from severe performance degradation.

Little attention has been paid in published reports on SAS to water-column influences. This is due to the fact that all current SAS systems are relatively *short-range* systems. The maximum detection range that can be achieved with standard SAS depends on vehicle speed V and array aperture L. A new ping must be transmitted before half the array aperture is travelled [10] in order to avoid spatial undersampling. Hence the minimum pulse repetition frequency $PRF = 2V/L$ from which the maximum range R_{max} follows:

$$R_{max} = C / 2PRF = C\,L / 4V,$$

where C is the speed of sound ($\approx$ 1500 m/s). For typical SAS systems $V = 2$ m/s and $L = 1$ m, such that $R_{max} \approx 190$ m. At these short ranges the devastating effects from the water column have had no time to appear. However, in the SAS project we aim at longer ranges ($R_{max} \approx 500$ m and more) in order to increase the coverage rate (see Section 4.7). Influences of the water column may become significant here, which opens up a new area of research.

4.3 Surface reverberation

Surface reverberation is more complicated than sea floor reverberation, due to the fact that it varies over time. This makes the subject rather untouchable, and it is not difficult to understand why other authors do not mention the problem. However, the influence of surface reflections can be minimised by using vertical directivity, combined with proper shading of both transmitter and receiver elements to suppress side lobes. An example of hardware shading of a transmitter element,

designed by DERA, is shown in Figure 4.1. Remaining surface reflections will tend to smear in the SAS image, owing to their variability over time. We can therefore conclude that the degrading effect of surface reverberation is not likely to be a dominant factor and can be avoided to a large extent by proper system design.

Figure 4.1: Transmit array element with mechanical shading.

4.4 Multi-path in shallow water

It is rather obvious that in shallow water multi-path arrivals on the receiver will cause interference. The wide bandwidth provides enough range resolution to separate most multi-paths for a single highlight, but cannot help the destruction of the highlight structure of a target. This will seriously hinder highlight-based classification methods. Moreover, multi-path propagation tends to fill target shadows, which also reduces classification performance. Techniques to compensate for multi-paths (adaptive and/or model based matched filters) are currently being developed at DERA and TNO Defence, Security and Safety (and many other places), but these are far from mature.

Another consequence of multi-path is degradation of autofocussing performance, which is caused by decreasing correlation values as a result of multi-path. No published reports have been found on solutions for this problem. Possibly multi-path suppression techniques used in mobile communication could also be applied in SAS.

4.5 Array motion compensation by autofocussing

If a UUV is used in SAS, the array orientation (heading, roll and pitch) will vary over time. Since relatively long integration times are required to obtain sufficient resolution in the SAS process, array positioning will vary substantially in this time. Straight-forward synthetic aperture formation is not possible because of this array motion and some kind of motion compensation is required. Motion sensors are often used for this purpose in SAR, but wavelengths are very short in SAS (in the order of centimetres) and array positions are required to be one order of magnitude more accurate. Positioning sensors (heading, roll and pitch, accelerometers) are generally not sufficiently accurate. Therefore, autofocussing methods are widely applied in SAS [15 - [23]. The following list of methods originates from the report by Shippey and colleagues [15] and has been updated according to the latest points of view as presented at the ECUA 2000 conference, in which auto-focussing in SAS was a big issue.

Power spectra estimation uses the received signal from a 'target of opportunity'. The phase of reflected signals will be corrupted by the positioning error of the array. This measured phase could replace the phase of the transmitted signal in the matched filter.

Problems with this method are manifold. A proper point-like target of opportunity should be present, and have a sufficient SNR such that the phase can be measured before matched filtering. A bottom feature often has a too complex structure and the phase law of such a reflection is a sum of multiple arrivals, which disturbs the picture dramatically.

Furthermore, many other effects than the array positioning error, e.g., water-column influences, disturb the phase of the received signal, often even to a larger extent than the array positioning error.

Multi-look registration is a method in which the aperture is divided into sub-apertures. These are beam-formed and the position of the sub-apertures relative to a target of opportunity is determined by means of cross-correlation. The positioning of the whole aperture can be determined from the relative positions of the sub-apertures. "Aperture" is mentioned above and not "array" because the game can be played on the real aperture as well as on the synthetic aperture.

Yet this method is not very popular. If the real aperture is divided, the method lacks accuracy and if the synthetic aperture is used, some of the desired azimuth resolution is lost by the use of the smaller apertures. The latter is unacceptable for most SAS systems.

An option that has not yet been studied in this respect is the combination of beam-space adaptive beam-forming (as in Liu's thesis [9]) and multi-look autofocussing.

Beam-space adaptive beam-forming also uses sub-apertures; combining these methods may therefore be very cost-effective.

Phase gradient autofocus (PGA) is a very popular method in SAR [3] and has also been recommended for SAS by several authors, such as Gough and colleagues [17 - [19]. A target of opportunity of good SNR is also required here. The target is followed and its track (before SAS) is assumed to be a parabola; cf Figure 2-3. Deviations of this parabola are assumed to be caused by the array motion and can be compensated for (the exact form of the function is a hyperbola, but if the synthetic aperture is not too large, this can be adequately approximated by a parabola). The phases are adjusted by looking at the expected shape of the track, i.e., by looking at the phase gradient. Successes have been reported using this method.

Correlation methods are also popular. In these methods, the cross-correlation coefficient of successive ping outputs is maximised as a function of the motion parameters (e.g., speed, heading). It is assumed that two successive sonar images are highly correlated, which is generally true if many fixed scatterers (bottom reverberation) are present and the PRF is high (this is the link to the under--sampling problem). Several varieties of this method exist. The patent on the algorithm of displaced phase centres (DPC) is already old [22], but this algorithm is still hard to beat. It is implemented in the SAS processor of GESMA and successes on the rail data (available to TNO Defence, Security and Safety) have been obtained. Saclantcen also uses this method [20] in their processing. A similar method is the ping-to-ping cross-correlation method (P2C2) [23].

Contrast optimisation is recommended by Griffiths and colleagues [15]. This method maximises the contrast (signal-to-background-noise ratio) of a target of opportunity as a function of one motion parameter. Successes have been reported in laboratory experiments by DERA [16]. The method makes sense since the outcome of the autofocussing is exactly what you want, i.e., more contrast. Other methods come up with phase repairs which only indirectly enhance contrast. Whether this will work for UUV operations is another point. In that case not one, but four parameters (speed, heading, roll and pitch) need to be estimated. This is computationally and mathematically much more difficult; very often local, rather than global optimisation is found.

Note that all five above-mentioned methods depend on the presence of a high SNR target of opportunity. Preferably this is a point-like target, especially for the first three methods, which rely on phase 'repairing'. Very often only bottom reverberation is present for the autofocussing. It should be studied whether the above algorithms are sufficiently robust to use bottom reverberation for autofocussing. This is an interesting task, which clearly interacts with the topic of Section 4.1 above. Probably the best way to obtain a robust method for motion compensation is to combine motion sensors with autofocussing.

4.6 Interferometry

Interferometry is a technique not restricted to SAS in which phase differences between the signals from two vertically spaced receivers are used to derive height information. When applied to SAS, the technique requires the comparison between two SAS images with properly reconstructed phases. These phases should be known without the usual 2π ambiguity. Deriving the total phase from the original 'wrapped' phases is also known as *phase unwrapping*, which is the bottleneck in deriving interferometric images.

Interferometric SAS is discussed by various authors [12 - [15], and phase unwrapping is emphasized by Banks and colleagues [13]. Perrot and colleagues [14] and Griffiths and colleagues [15] concentrate on high-resolution 3-D seabed mapping.

4.7 Increased coverage rate

Two things are required in order to obtain a higher coverage rate: high search speed and long detection range. The coverage rate (in m^2/s) is given by: $V R_{max}$. There is a conflict if high resolution (by using SAS) is also required. The maximum range in standard SAS techniques is given by $R_{max} = C L / 4V$ as mentioned before in Section 4.2. This means that, obeying the spatial sampling criterion, the maximum coverage rate for a SAS system is fixed at: $C L / 4$.

This only depends on the array length, so that once this is fixed, nothing further can be done to improve this. It has been stated earlier that an important ambition of the SAS project is to improve the coverage rate nonetheless. This will be a challenging task.

The problem underlying the low coverage rate of SAS is the so-called *under-sampling problem* [4, [5, [15,[19]. In order to create a sufficiently sampled synthetic aperture, the associated real apertures should be at least half overlapping [10]. If the sampling is less dense than this, the synthetic array has holes in its aperture, which leads to a reduced beam pattern quality with higher side lobes and the possible formation of grating lobes. Grating lobes are spatial aliases, which manifest as ghost targets in the SAS image.

It is clear that this is unwanted. Therefore all current SAS systems sample correctly, accepting the low coverage rate. In doing so, they also help autofocussing, since most autofocus algorithms rely on sufficiently high sampling rates (more than half overlapping arrays). The under-sampling problem is therefore also related to the autofocussing problem of Section 4.4.

Almost all papers on SAS mention the problem and a lot of effort has been put into solving the under-sampling problem. Some authors have suggested solutions And we shall follow Griffiths and colleagues [15] in which most ideas are summarised:

Vehicle speed reduction: A low vehicle speed helps the under-sampling problem and/or allows for a longer range, but it does not create a higher coverage rate. Furthermore, vehicle stability requires a minimum speed, so that this suggestion is not useful.

Wideband: Wideband has some favourable properties for grating lobes. The location of grating lobes is frequency dependent: high-frequency ghost images are located closer to the real image than low–frequency images. The grating lobes tend to smear when the frequency band is wide. Some authors claim that by using wideband we can get away with serious under-sampling and that the PRF can be reduced without suffering from ghost images. This looks very promising; however, autofocussing was not applied for motion compensation in any of the quoted studies (they were either theoretical studies or rail experiments). Thus, although the aliasing problem may not seem so severe, wideband does not solve the array positioning. On the contrary, the larger the bandwidth (= range resolution) the more sensitive the processor is to positioning errors.

The only real solution seems to be having multiple pings in the water simultaneously. Some proposed ways to achieve this are reviewed below.

Orthogonal pulses: Orthogonal pulses are also a promising way to increase the vehicle speed without lowering the PRF. Orthogonal pulses are pulses that have minimal correlation and therefore give low output in each others matched filter. Well-known examples are up-sweep and down-sweep FM. These pulses exhibit minimal interference and are in the same frequency band. If alternately transmitted, the matched filter of the down-sweep will suppress the energy of the up-sweep of the previous ping, which can now be used for twice the maximum range. Pulses exist with double phases (biphase codes) which have two sets of orthogonal pulses. These pulses have successfully been tested in a SAR application [12]. This would increase the maximum SAS range by a factor of 4, bringing it close to the ambition of the SAS programme. Another factor of 2 could be gained by splitting the bandwidth (if sufficiently wide, this is an option).

Orthogonal pulses seem the ultimate solution for the under-sampling problem in SAS, but there are also problems or restrictions:

- The first is practical: the receiver elements and amplifier electronics should not saturate during signal transmission, since echoes of the previous ping will continue to arrive during transmission of the next. This is a severe demand on the dynamic range of the elements and the subsequent data acquisition system.
- The second is more fundamental: are orthogonal pulses sufficiently orthogonal, i.e., are all reverberations from a previous pulse sufficiently suppressed by the

matched filter? It could be argued that for wideband pulses this is generally the case, at least for point-like scatterers. However, is this true for extended scatterers, such as the sea floor, and how wide is sufficiently wide? A quantitative study on this subject is presented in a TNO Defence, Security and Safety technical note [35].

Multiple vertical beams: We can make vertical beams for each range, in a half-overlapping way by making a transmitter with vertical directivity. This enables autofocussing. The under-sampling problem (grating lobes) can be solved by wideband. This solution looks very simple and straight forward, but some problems can also be foreseen here.

The first problem is practical: a small transducer with a high horizontal and vertical directivity is required and these are conflicting requirements. The second is geometric: at long ranges in shallow water, beams must be extremely narrow and accurately steered, which puts high demands on vehicle stability and/or beam stabilization. However because of its simplicity, this solution is very attractive and worthy of further study.

Kiwi-SAS: There seems to be a trade-off in all these solutions between solving the under-sampling problem and minimising positioning errors. The only way really to improve the coverage rate is to relate the two problems, and tackle them at the same time. A first attempt at such an approach was published by the New Zealand researchers Gough and Hawkins [6, [19, [46]. Their solution, Kiwi-SAS, is worth studying.

4.8 Buried mine detection

For buried mine detection, low frequencies are required that can penetrate the sediment. Standard sonar techniques would require an impractically large receive array to maintain sufficient azimuth resolution. High azimuth resolution and low frequency can be combined to achieve bottom penetrating sonar images by using synthetic aperture techniques. No published reports have been found on this subject, but oral communication with AFN partners suggests that GESMA and Saclantcen have been experimenting successfully with low frequency, buried-mine SAS.

4.9 Shadow and echo classification

Current mine classification techniques rely on the acoustic shadow of a target, which indicates the size and shape of that object. Owing to the high resolution of SAS, more accurate shadows can be detected at longer ranges, thereby improving the possibilities for classification. The *shape* of a shadow, for instance, can be

detected using SAS and provides information about the shape of the target. Malkasse [48] describes an interesting technique in which multi-aspect shadows are combined to derive the original shape of the object using an inverse ray-tracing technique.

5. Discussion

The technique of SAS is very suitable for obtaining high-resolution images of the sea floor or sea mines. Expected operational benefits compared to conventional techniques are increased SNR (detection), sharper shadows by contrast enhancement (classification) and increased resolution (localisation).

The technique is very similar to SAR in theory and should be easily applicable. In practice, however, there are several complications, most of which have been studied in the TNO Defence, Security and Safety SAS project (see Chapter 4). The main problems are limited coverage rate and array motion compensation. These problems need proper consideration and the number of reported studies on these topics is growing rapidly.

SAS exists in theory, and it is also being experimentally proven; by DERA in the laboratory [15] and in rail experiments with their transportable rail facility, at GESMA in rail experiments with the Lanvéoc rail facility [44, [45] and even with a UUV in New Zealand [46] and Sweden [47]. The analysis results of rail data in the TNO Defence, Security and Safety SAS project also look promising. Two examples of SAS images from our own analysis can be seen in Figure 5.1 and Figure 5.2.

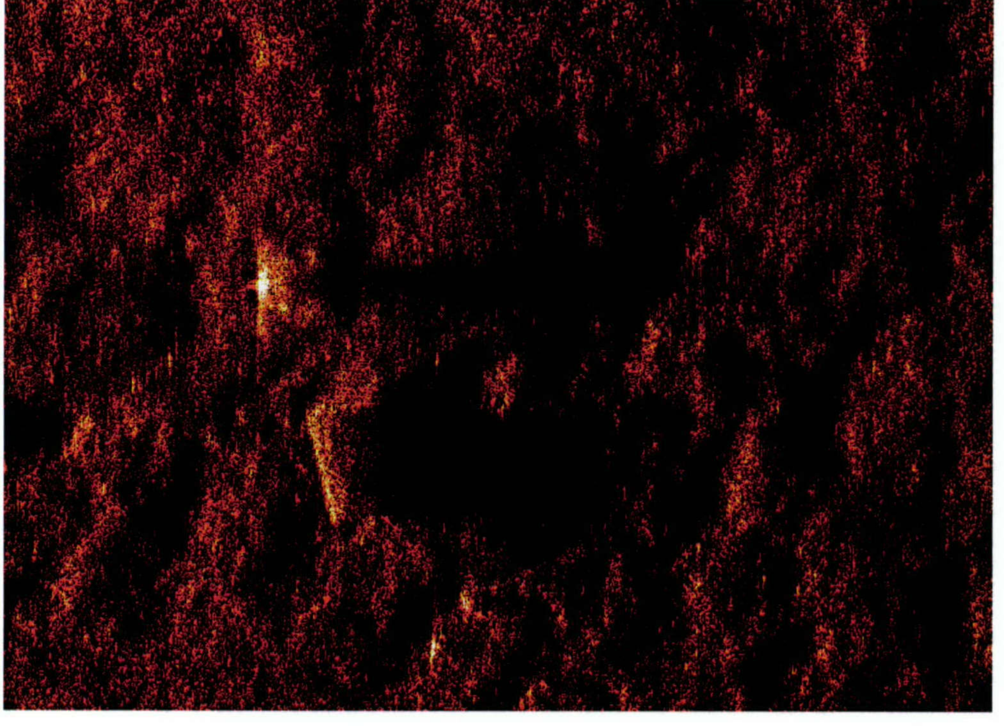

Figure 5.1: SAS image of a sphere (upper target) and a cylinder (lower target) at 75 m range from a 10 m long rail.

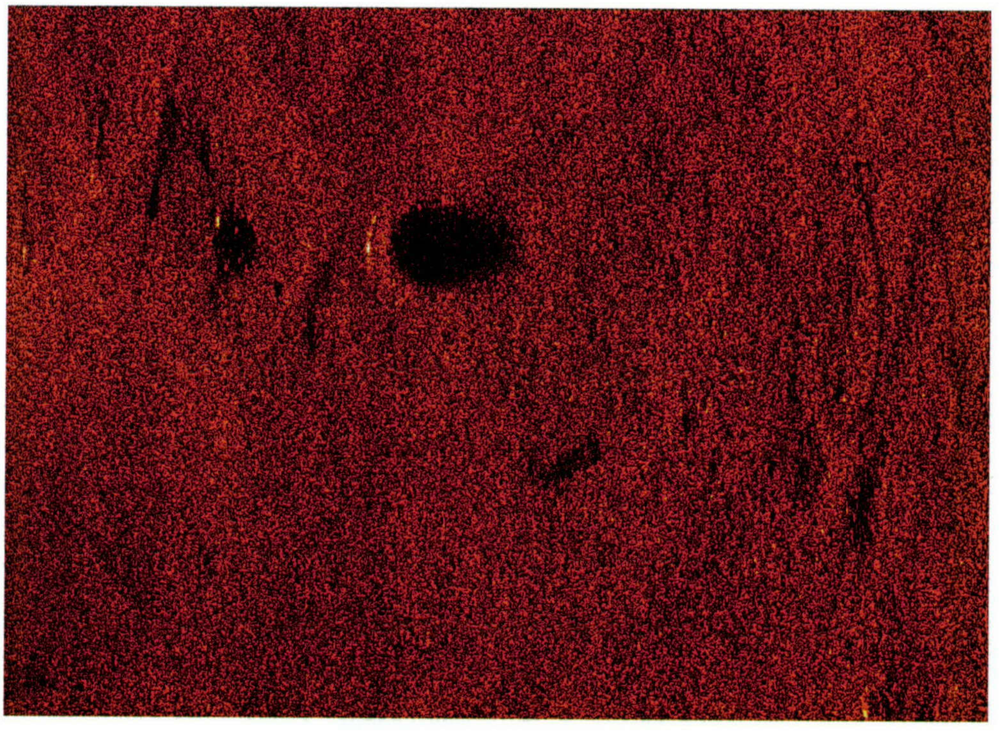

Figure 5.2: SAS image of a sphere at 25 m range from a 10 m long rail.

These images have been obtained from experimental data collected during a rail trial in 1999 by our SAS-project partners at DERA and GESMA. Targets were placed at 25, 50 and 75 m from the 10 m long rail and echoes from the 150 kHz, 60 kHz bandwidth sonar were recorded for off-line analysis. Each recording contained 320 pings (LFM chirps of 4 ms duration), received by two arrays of 32 elements each. The images above were obtained from only one of the two arrays. The sonar images resemble a photograph: target shadows give a clear indication of the target shape, and the bottom structure can also be observed.

The published reports found and the analysis results presented indicate that SAS seems to fulfil its promises. In the future, further maturation of the technique is to be expected towards more operational settings and on realistic platforms. The work at TNO Defence, Security and Safety will be aimed at improving processing (efficiency, accuracy), auto-focussing and motion compensation.

6. References

Review of SAR

[1] J.C. Curlander. *Synthetic Aperture Radar*, Wiley, 1991.

[2] W.G. Carrara, R.S. Goodman & R.M. Majewski. *Spotlight synthetic aperture radar*. Artech house, 1995.

[3] M.P.G. Otten. *Review of SAR processing techniques*. European MW 1998.

Review of SAS

[4] L.J. Cutrona: *Comparison of sonar system performance achievable using synthetic aperture techniques with the performance achievable by more conventional means*, , J. Acoust. Soc. Am, **58** (2), pp. 336-348, 1975.

[5] L.J. Cutrona: *Additional characteristics of synthetic aperture sonar systems, and a further comparison with nonsynthetic aperture sonar systems,* J. Acoust. Soc. Am, **61** (2), pp. 1213-1217, 1977.

[6] P.T. Gough & D.W. Hawkins, *A Short History of Synthetic Aperture Sonar,* IEEE, 1998.

[7] P. de Heering, *Acoustic synthetic aperture processing: theory and application*, PhD thesis, University of Bremen, Germany, 1990.

[8] D.W. Hawkins, *Synthetic Aperture Imaging Algorithms: with application to wide bandwidth sonar*, PhD thesis, University of Canterbury, Christchurch, New Zealand, 1996.

[9] Qianjun Liu, *Adaptive Imaging Using Wideband Active Sonar,* PhD thesis, Chambers University of Technology, Göteborg, Sweden, 1999.

SAS beamforming

[10] J. Châtillon, *Absence of Nyquist criterion for spatial sampling of a strip map Synthetic Aperture Sonar,* Proceedings of ECUA 2000, pp. 401-406

[11] H.D. Griffiths et. al., *Adaptive SAR beamforming network.,* Report ESTEC 6553/89/NL/IW, Noordwijk, 1992.

[12] H.D. Griffiths et. al., *Mine detection using interferometric Synthetic Aperture Sonar,* Proceedings of UDT, 1995.

[13] S. Banks, T. Sutton & H. Griffiths. *Noise susceptibility of phase unwrapping algorithms for interferometric synthetic aperture sonar*. Proc. ECUA2000, pp. 451-456

[14] Y. Perrot, B. Hamonic & M. Legris. *Three-dimensional high resolution imaging using SAS and interferometric process.* Proc. ECUA2000, pp. 425-432.

Autofocussing

[15] H.D. Griffiths et. al., *Interferometric synthetic aperture sonar for high-resolution 3D mapping of the seabed.,* IEE Proc. Radar, Sonar Navig., **144** (2), 1997.

[16] T.J. Sutton, S.A. Chapman & H.D. Griffiths, *Robustness and effectiveness of autofocus algorithms applied to diverse seabed environments.* Proc. of ECUA 2000, pp. 407-412.

[17] K.A. Johnson, M.P. Hayes & P.T. Gough: *A method of estimating the sub-wavelength sway of a sonar tow fish,* IEEE Journal of Oceanic Engineering, **20** (4), pp. 258-267, 1995.

[18] P.T. Gough & D.W. Hawkins, *Imaging Algorithms for a Strip-Map Synthetic Aperture Sonar: Minimizing the Effects of Aperture Errors and Aperture Undersampling,* IEEE Journal of Oceanic Eng., **22** (1), 1997.

[19] P.T. Gough, M.P. Hayes & H.J. Callow, *Strip-Map Path correction using phase matching autofocus,* Proceedings of ECUA 2000, pp. 413-418.

[20] M. Pinto, S. Fioravanti & E. Bovio, *Accuracy of synthetic aperture sonar micronavigation using a displaced phase center antenna.* SACLANTCEN Memorandum SM-352, February 1999.

[21] A. Belletini & M. Pinto. *Experimental investigation of synthetic aperture sonar micronavigation.* Proc. ECUA2000, pp. 445-450.

[22] R.S. Raven, *Electronic stabilization for displaced phase center,* Patent 4.244.036, 1978.

[23] D. Billon, *Interferometric synthetic aperture sonar design and performance issues,* Proceedings of ECUA 2000

Computational load reduction

[24] I. Ekhaus: *Reduced wavenumber synthetic aperture for MCM applications,* Technical Report, Synap Corporation, 1997

[25] G. Shippey, P. Ulriksen & Qianjun Liu: *Quasi narrowband processing of wideband sonar echoes,* Proc. of the Fourth European Conference on Underwater Acoustics, 1998.

[26] G. Shippey, Qianjun Liu & P. Ulriksen: *Wideband synthetic aperture sonar processing using a time-domain approach,* IEEE Journal of Oceanic Engineering, 1999.

[27] L.M.H. Ulander, H. Hellsten & G. Stenström: *Synthetic aperture radar processing using Fast Factorised Backprojection,* Proc ??, 1999.

[28] L.C. Godara & M.R.S. Jahromi: *Limitations and capabilities of frequency domain broadband constrained beamforming schemes,* IEEE **47** (9), pp. 2386-2393,1999.

[29] P.T. Gough, M.P. Hayes & D.R. Wilkinson, *An efficient image reconstruction algorithm for a multiple hydrophone array synthetic aperture sonar* , Proceedings of ECUA 2000.

Sonar array beamforming

[30] R.O. Nielsen, *Sonar Signal Processing, Artech House, 1991.*

[31] S. Haykin, J.H. Justice, N.L. Owsley, J.L. Yen & A.C. Kak, *Array Signal Processing, Prentice-Hall,* 1985.

[32] K. Jones, *ROSCO, a robust beamforming algorithm*, DERA 1995.

[33] P.A.M de Theije, *Auto-calibration methods for towed arrays,* TNO-FEL report in progress.

[34] S.P. Beerens, R. Been, J. Groen, Y. Doisy & E. Noutary, *Adaptive PS beamforming for triplet array sonar*, submitted to IEEE Oceanic Engineering 2000.

[35] S.P. Beerens, *The use of orthogonal pulses in active sonar applications*, Technical note, TNO-FEL, 2000

[36] R. Been, S.P. Beerens & C. Volwerk, *Final report of the Shape Corrected Beamforming SCB programme*, TNO-FEL report FEL-97-A384, 1998.

Passive SAS

[37] S. Stergiopoulos: *Implementation of Adaptive and Synthetic-Aperture Processing, Schemes in Integrated Active-Passive Sonar Systems,* Proc. IEEE, **86** (2), pp 356-396, 1998.

[38] S. Stergiopoulos & E.J. Sullivan: *Extended Towed Array processing by an overlap correlator*, J. Acoustic. Soc. Am, **86**, pp. 158-171,1989.

[39] S. Stergiopoulos: *Optimum bearing resolution for a moving Towed Array and extension of its physical aperture*, J. Acoustic. Soc. Am, **87**, pp. 2128-2140, 1990.

[40] S. Stergiopoulos: *Limitations on Towed Array gain imposed by a non-isotropic ocean,* J. Acoustic. Soc. Am, **88**, 1991.

[41] S. Stergiopoulos & H. Urban: *A new Passive Synthetic Aperture Technique for Towed Arrays,* IEEE Journal of Oceanic Engineering, **17** (1), pp. 17-25, 1992.

[42] R.E. Williams: *Creating an acoustic synthetic aperture in the ocean,* J. Acoustic. Soc. Am, **60** (4), pp. 60-73., 1976

[43] R.E. Williams & B. Harris: *Passive acoustic synthetic aperture processing techniques,* IEEE Journal of Oceanic Engineering, **17** (1), pp. 8-16, 1992

Experimental SAS

[44] S. Guyonic, A. Salaün & C. Queïnnec, *Low and medium frequency synthetic aperture sonar experiments on a rail,* GESMA, 1997.

[45] A. Hétet. *Evaluation of specific aspects of synthetic aperture sonar, by conducting at sea experiments with a rail, in the frame of mine hunting systems design.* Proc. of ECUA2000, pp. 439-444

[46] P.T. Gough & D.W. Hawkins: *Test results using a prototype synthetic aperture sonar,* Journal of the Acoustic Society of America, **86** (6), pp. 2328-2333, 1997.

[47] G. Shippey, Presentation at TNO-FEL, 18-9-2000.

Shadow classification

[48] J.P. Malkasse. *Identification tool by sonar means with synthetic aperture processing and 3D reconstruction.* Proc. UDT Europe 2000, session 7B.1.

[49] P. Galerne, e.a. *Object classification using neural networks in sonar imagery.* SPIE proc. Vol. 3101, 1997. ISBN 0-8194-2521-4, pp. 306-314.

7. Signature

F.P.G. Driessen
Group leader

J.C. Sabel
Author

S.P. Beerens
Author

J. Groen
Author

REPORT DOCUMENTATION PAGE
(MOD-NL)

1. DEFENCE REPORT NO (MOD-NL)	**2. RECIPIENT'S ACCESSION NO**	**3. PERFORMING ORGANIZATION REPORT NO**
TD04-0415		TNO-DV1 2004 A253
4. PROJECT/TASK/WORK UNIT NO	**5. CONTRACT NO**	**6. REPORT DATE**
015.33053	A99KM757	April 2005
7. NUMBER OF PAGES	**8. NUMBER OF REFERENCES**	**9. TYPE OF REPORT AND DATES COVERED**
37 (incl appendices, excl RDP & distribution list)	49	Final

10. TITLE AND SUBTITLE

A literature survey on Synthetic Aperture Sonar (SAS)

11. AUTHOR(S)

Dr.Ir. J.C. Sabel, Dr. S.P. Beerens and Ir. J.Groen

12. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

TNO Physics and Electronics Laboratory, PO Box 96864, 2509 JG The Hague, The Netherlands
Oude Waalsdorperweg 63, The Hague, The Netherlands

13. SPONSORING AGENCY NAME(S) AND ADDRESS(ES)

DMKM/WCS/OSWS

14. SUPPLEMENTARY NOTES

The classification designation Ongerubriceerd is equivalent to Unclassified, Stg. Confidentieel is equivalent to Confidential and Stg. Geheim is equivalent to Secret.

15. ABSTRACT (MAXIMUM 200 WORDS (1044 BYTE))

Synthetic Aperture Sonar (SAS) is a promising new technique for application in small Unmanned Underwater Vehicles (UUVs) for future mine-hunting operations. The relatively short array in a UUV can be synthetically extended to obtain higher resolution and more array gain, leading to improved detection, classification and localisation of mines. Unlike its radar equivalent Synthetic Aperture Radar (SAR) the SAS technique is not yet mature. On the contrary, SAS is still very much unproven, while SAR has already been operational for many years. As a consequence, literature on SAR is abundant, but on SAS relatively rare. This report gives an overview of relevant literature, in which the most fundamental problems of SAS are reviewed.

16. DESCRIPTORS	**IDENTIFIERS**
Synthetic aperture sonar	

17a. SECURITY CLASSIFICATION (OF REPORT)	**17b. SECURITY CLASSIFICATION (OF PAGE)**	**17c. SECURITY CLASSIFICATION (OF ABSTRACT)**
Ongerubriceerd	Ongerubriceerd	Ongerubriceerd

18. DISTRIBUTION AVAILABILITY STATEMENT	**17d. SECURITY CLASSIFICATION (OF TITLES)**
Unlimited	Ongerubriceerd

Distributielijst

1. SC-WOO
2. HWO-KM
3. DMKM/WCS/OSWS, t.a.v. LTZE 2OC J.P. Gillesse
4. Bibliotheek KMA
5. Bibliotheek KMA
6. Bibliotheek KMA
7. Directeur Operaties TNO Defensie en Veiligheid
8. Directeur Kennis TNO Defensie en Veiligheid, daarna reserve
9. Archief TNO Defensie en Veiligheid, in bruikleen aan Accountmanager KM
10. Archief TNO Defensie en Veiligheid, in bruikleen aan drs. C.M. Ort
11. Archief TNO Defensie en Veiligheid, in bruikleen aan dr. S.P. Beerens
12. Archief TNO Defensie en Veiligheid, in bruikleen aan ir. F.P.G. Driessen
13. Archief TNO Defensie en Veiligheid, in bruikleen aan ir. J. Groen
14. Archief TNO Defensie en Veiligheid, in bruikleen aan dr.ir. J.C. Sabel
15. Archief TNO Defensie en Veiligheid, in bruikleen aan ir. P. Schippers
16. Documentatie TNO Defensie en Veiligheid
17. Reserve

Distributie managementuittreksel & distributielijst

1×	Directeur Markt TNO Defensie en Veiligheid, Algemeen Directeur TNO Defensie en Veiligheid, MIVD / AAR / HBMT
4×	SC-WOO, HWO-KM, HWO-KLu, HWO-CO
8×	HWO-KL

Indien binnen de krijgsmacht extra exemplaren van dit rapport worden gewenst door personen of instanties die niet op de verzendlijst voorkomen, dan dienen deze aangevraagd te worden bij het betreffende Hoofd Wetenschappelijk Onderzoek of, indien het een K-opdracht betreft, bij de Directeur Wetenschappelijk Onderzoek en Ontwikkeling.